

Attorney Docket No. 16869K-090500US
Client Ref. No. 630/SM/co

PATENT APPLICATION

Storage System, Storage System Control Method, and Storage Medium Having Program Recorded Thereon

Inventors: **Akihiro MORI**
Citizenship: Japan

Assignee: Hitachi, Ltd.
6, Kanda Surugadai 4-chome
Chiyoda-ku, Tokyo, Japan
Incorporation: Japan

Entity: Large

TOWNSEND AND TOWNSEND and CREW LLP
Two Embarcadero Center, 8th Floor
San Francisco, California 94111-3834
(415) 576-0200

STORAGE SYSTEM, STORAGE SYSTEM CONTROL METHOD, AND
STORAGE MEDIUM HAVING PROGRAM RECORDED THEREON

CROSS-REFERENCE TO RELATED APPLICATIONS

5 The present application claims priority upon Japanese Patent Application No. 2002-300789 filed October 15, 2002, which is herein incorporated by reference.

BACKGROUND OF THE INVENTION

10 Field of the Invention

 The present invention relates to a storage system, a method of controlling a storage system, and a storage medium having a control program recorded thereon.

Description of Related Art

15 In a storage system including a host computer and a disk array device connected with each other so that they can communicate, there is known a duplication management function which enables duplication of data among logical volumes according to real-time copying.

20 The duplication management function will be described with reference to Fig. 11A through Fig. 11D. In Fig. 11A, a logical volume 71 ("primary logical volume" hereinbelow) formed on physical storage areas provided by disk drives 25, which serve as storage devices, is a logical volume that can be directly
25 accessed from an operating system 12 ("OS" hereinbelow) running on a host computer 10.

 On the other hand, another logical volume 72 ("secondary logical volume" hereinbelow) stores a duplicate of the data stored in the primary logical volume 71 by means of the duplication
30 management function of the storage system.

 In response to, for example, an instruction from the user,

the duplication management function makes the control state of the primary logical volume 71 and the secondary logical volume 72 change between the "paired state" and the "split state". In the "paired state", when the contents of the primary logical volume 71 is updated, the contents of the secondary logical volume 72 will also be updated in a real-time basis so as to ensure consistency between the contents of the primary logical volume 71 and that of the secondary logical volume 72. Note that in the "paired state", access from another host computer 100 other than the host computer 10 to the secondary logical volume 72 is prohibited.

The "paired state" is cancelled, for example, when a pair split command sent from the host computer 10 by the user's instruction is received by the disk array device 20, and according to this command, the control state of the primary logical volume 71 and the secondary logical volume 72 is changed to the "split state" (see Fig. 11B). In the "split state," the other host computer 100, which was prohibited to access the secondary logical volume 72, will be allowed to access this secondary volume 72. The other computer 100 will be able to access the secondary logical volume 72 during the "split state" in order to backup data, perform batch processing or data analysis, and so forth. Note that in the "split state," any update to the primary logical volume 71 from the host computer 10 will not be reflected to the secondary logical volume 72.

In usual operation, data that has not been reflected to the secondary logical volume 72 and that has been generated by updating in the primary logical volume 71 is stored as "differential data." This differential data is to be used to make the contents of the secondary volume 72 consistent with the primary volume 71 prior to returning to the "paired state", as will be described later.

When a command to return from the "split state" to the "paired state" (such a command is referred to as a "resync command" hereinbelow) is issued from the host computer 10 to the disk array device 20, the duplication management function will synchronize the contents of both volumes 71, 72 (resync in progress) using the differential data, and then cause transition of the control state of primary logical volume 71 and secondary logical volume 72 to the "paired state".

As described above, the secondary logical volume 72 in the "split state" may be accessed from an OS 102 running on the host computer 100, and the computer 100 may make use of data in the secondary logical volume 72 to backup data, perform any batch operation or data analysis, and so forth.

During the "split state", however, the host computer 10, which accesses the primary logical volume 71, may not have access to the secondary logical volume 72 as an independent volume separate from the primary logical volume 71. This is because the OS 12 running on the host computer 10 stores identical logical volume names and data set names in the management information for the primary logical volume 71 and the secondary logical volume 72 that is managed by the duplication management function.

The management information may include VTOC (Volume Table of Contents), catalog information, VSAM (Virtual Storage Access Method) data set, and VVDS (Vsam Volume Data Set) that is the information about the data set of the target of system management, if the OS 12 is one for a mainframe. The management information may be the file system management information if the OS 12 is one for an open system.

From the point of view of, for example, attempting to increase performance and decrease cost for a host computer 10, there exists a need to make the secondary logical volume 72

accessible as a logical volume independent of the primary logical volume 71 from the host computer 10, which manages both the primary and secondary logical volumes 71, 72, and not only from the other computer 100.

5 In order to make the secondary logical volume 72 accessible as a logical volume independent of the primary logical volume 71, it may be sufficient to configure the logical volume names and data set names described in the management information for the primary logical volume 71 and management information for the
10 secondary logical volume 72 to be different from each other.

Such a technology is known in the art and disclosed, for example, in Japanese Patent Application Laid-open Publication No. 2000-339104.

15 However, in general, the volume names and data set names are described in a number of locations in the management information. It is very cumbersome to perform settings for each data by using an editor, a tool, etc., and in addition, a human error might occur during the setting. It is not practical to perform setting for volume identifiers and data set identifiers
20 every time a routine job, such as a daily data backup, a daily batch processing, and data analysis, is performed at a regular or irregular interval.

BRIEF SUMMARY OF THE INVENTION

25 The present invention has been made in view of the above circumstances and has an object to provide a method of controlling a storage system, a storage system, and a storage medium having a control program recorded thereon.

30 An aspect of the storage system of the present invention for achieving the above object is as follows:

The storage system has a storage control device connected

to a host computer to be able to communicate therewith for controlling the data input/output (I/O) to/from a storage device according to a data I/O request sent from the host computer.

5 The storage control device manages a storage area in the storage device using a logical volume that is a logical storage area created on the storage area.

The system stores, in the logical volume, management information that enables an operating system running on the host computer to manage the logical volume.

10 The storage control device controls duplication of data in a first logical volume also to be stored on a real-time basis (in synchronism) in a second logical volume that is different from the first logical volume.

15 During synchronized (real-time) duplication, a logical volume identifier of the first logical volume and a logical volume identifier of the second logical volume described in the management information are made to be consistent with (are made to match) each other, and a data set identifier of the first logical volume and a data set identifier of the second logical volume
20 described in the management information are made to be consistent with (are made to match) each other.

In the method of the present invention, the storage system generates a control program for performing a process for setting the logical volume identifier and the data set identifier for the
25 first logical volume, which are described in the management information in the first logical volume, and the logical volume identifier and the data set identifier for the second logical volume, which are described in the management information in the second logical volume, to be different from each other.

30 Then, the storage system executes the control program after interrupting the duplication in order to make the second logical

volume be recognized as being accessible by the operating system as a logical volume independent of the first logical volume.

Features and objects of the present invention other than the above will become clear by reading the description of the present specification with reference to the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

For complete understanding of the present invention and the advantages thereof, reference is now made to the following description taken in conjunction with the accompanying drawings in which:

Fig. 1 is diagram illustrating a schematic configuration of a storage system according to an embodiment of the present invention;

Fig. 2 is a flow chart illustrating a first process flow according to the embodiment of the present invention;

Fig. 3 is a flow chart illustrating a second process flow according to the embodiment of the present invention;

Fig. 4 is a schematic diagram illustrating states of the volume name, the data set name, and the catalog name when the secondary logical volume is recorded in a catalog different from that of the primary logical volume, in accordance with the embodiment of the present invention;

Fig. 5 is a schematic diagram illustrating states of the volume name, the data set name, and the catalog name when the secondary logical volume is recorded in the same catalog as that of the primary logical volume, in accordance with the embodiment of the present invention;

Fig. 6 is a schematic diagram illustrating states of the volume name, the data set name, and the catalog name when the

secondary logical volume is recorded both in the same catalogs as those of the primary logical volume and in catalogs different from those of the primary logical volume, in accordance with the embodiment of the present invention;

5 Fig. 7 is a diagram showing an example of a control program when the secondary logical volume is to be recorded in a catalog different from that of the primary logical volume, in accordance with the embodiment of the present invention;

10 Fig. 8 is a diagram showing an example of a control program when the secondary logical volume is to be recorded in the same catalog as that of the primary logical volume, in accordance with the embodiment of the present invention;

15 Fig. 9 is a diagram showing an example of a control program when the secondary logical volume is to be recorded both in the same catalogs as those of the primary logical volume and in catalogs different from those of the primary logical volume, in accordance with the embodiment of the present invention;

 Fig. 10 is a flow chart illustrating a third process flow according to the embodiment of the present invention; and

20 Fig. 11A through Fig. 11D are schematic diagrams illustrating a duplication management function of a storage system.

DETAILED DESCRIPTION OF THE INVENTION

25 --Overview--

 At least following should be best understood by reading the description and from the accompanying drawings.

 The present invention provides a method of controlling a storage system. The storage system includes: a host computer;
30 and a storage control device that is connected to the host computer to be able to communicate therewith and that is for

inputting/outputting data to/from a storage device based on a data input/output request sent from the host computer. The storage system: manages a storage area provided by the storage device using a logical volume that is a logical storage area created on the storage area; stores, in the logical volume, management information that enables an operating system running on the host computer to manage the logical volume; controls duplication of data in a first logical volume also to be stored on a real-time basis in a second logical volume that is different from the first logical volume; and makes a logical volume identifier and a data set identifier for the first logical volume that are described in the management information in the first logical volume and a logical volume identifier and a data set identifier for the second logical volume that are described in the management information in the second logical volume match each other while the real-time duplication is being performed. In the method, the storage system: generates a control program for performing a process for setting the logical volume identifier and the data set identifier for the first logical volume, which are described in the management information in the first logical volume, and the logical volume identifier and the data set identifier for the second logical volume, which are described in the management information in the second logical volume, to be different from each other; and executes the control program after interrupting the duplication in order to make the second logical volume be recognized as being accessible by the operating system as a logical volume independent of the first logical volume.

For example, the storage device may be a disk drive. The data input/output request may be a data write request and/or data read request. The logical volume identifier may be a logical volume name, a volume label, etc. The data set identifier may

be a data set name, a file name, etc. The operating system may be an OS either of a mainframe or of an open system. The management information may be VTOC and/or VVDS.

In a storage system which is required to make consistent
5 the logical volume identifier and the data set identifier described in the management information of the first logical volume with those of the second logical volume during real-time duplication, by generating a control program for performing a process for setting the logical volume identifier and the data
10 set identifier for the first logical volume, which are described in the management information in the first logical volume, and the logical volume identifier and the data set identifier for the second logical volume, which are described in the management information in the second logical volume, to be different from
15 each other, and by executing the control program, after interrupting the duplication, in order to make the second logical volume be recognized as being accessible by the operating system as a logical volume independent of the first logical volume, the user does not have to perform settings of the logical volume
20 identifiers, the data set identifiers, etc. described in the management information every time it is necessary to access the second logical volume as a logical volume independent of the primary logical volume. Therefore, it is easy to make the secondary logical volume be recognized as accessible from the
25 operating system as a logical volume independent of the primary logical volume. By the automatic operations performed by the control program, there will be fewer occurrences of human errors; therefore, the second logical volume may be readily and easily recognized as a volume independent of the first logical volume
30 for the direct access such as data I/O.

Note that it is possible to include, in the above-described

control program, a control program for interrupting the real-time duplication, and the interruption of the real-time duplication may be performed by executing this control program. By doing this, the second logical volume may be recognized more easily for direct
5 access, such as for data I/O, as a volume separate from the first logical volume.

In addition, the operating system may manage the logical volumes by recording them in a catalog, and, when the second logical volume is being made to be recognized by the operating
10 system as a logical volume independent of the first logical volume, the second logical volume may be recorded in a catalog different from that of the first logical volume or in the same catalog as that of the first logical volume. In this way, the second logical volume can be made to be recognized by the operating system as
15 a volume separate from the first logical volume in a flexible manner corresponding to the user needs or operation strategy of the storage control device.

Further, when the second logical volume is being made to be recognized by the operating system as a logical volume
20 independent of the first logical volume, the storage system may generate a control program for performing a process for setting the logical volume identifier and the data set identifier for the first logical volume, which are described in the management information in the first logical volume, and the logical volume
25 identifier and the data set identifier for the second logical volume, which are described in the management information in the second logical volume, to be the same as each other and for restarting the real-time duplication, and the storage system may restart the real-time duplication by executing this control
30 program. The real-time duplication may be readily and positively resumed when the second logical volume has been made to be

recognized by the operating system as a separate volume independent of the first logical volume.

In addition, the second logical volume may be provided by a storage control device connected to the storage control device, 5 to be able to communicate therewith, that provides the first logical volume. In this way, even in case in which a duplication of the data in the first logical volume provided by the storage control device installed at a local site is to be stored in the second logical volume of another storage control device, which 10 is different from the above-described storage control device and installed at a remote site at a distance from that storage control device, the second logical volume may be recognized as a separate logical volume independent of the first logical volume.

15 --Embodiments--

An embodiment of a storage system according to the invention is shown in Fig. 1. A host computer 10 and a disk array device 20 are connected with each other via communication means 50 such as ESCON (registered trademark), a LAN (local area network), and 20 a SAN (storage area network). The host computer 10 may be either a mainframe computer or an open system computer including personal computers and workstations. An OS 12 is running on the host computer 10.

The disk array device 20 has channel controllers 21, disk 25 controllers 22, a cache memory 23, a control memory 24, and disk drives 25. The channel controllers 21, the disk controllers 22, the cache memory 23, and the control memory 24 are connected with each other through a switching means 26 such as a crossbar switch so that they can communicate. The channel controller 21 has a 30 communication interface to the host computer 10 for controlling data transfer between the host computer 10 and the cache memory

23. The disk controller 22 controls data transfer between the cache memory 23 and the disk drive 25.

The cache memory 23 is accessible from both the channel controllers 21 and the disk controllers 22 and stores data
5 input/output to/from the disk drives 25. The control memory 24 stores various kinds of information required for the operation of the disk array device 20 including the control information of the cache memory 23 and the disk drives 25. The control memory 24 is accessible from both the channel controllers 21 and the disk
10 controllers 22.

The disk drives 25 may be controlled according to RAID (Redundant Array of Inexpensive [or Independent] Disks). The physical storage areas provided by the disk drives 25 are managed by using logical volumes 70, which are the logical storage areas
15 created on the physical storage areas. In addition, in the storage system according to this embodiment, the duplication management function is made to work by being incorporated in the OS 12 or by executing a program that runs on the OS 12. In the "paired state," the contents of the primary logical volume 71 and
20 the contents of the secondary logical volume 72 are controlled to ensure their consistency on a real-time basis, so that the contents of secondary logical volume 72 will be updated in real-time when the contents of primary logical volume 71 is updated. It should be noted that since the functionality of duplication
25 management function itself is the same as that explained in the foregoing description, the description thereof is omitted here.

< Description of processing >

Now the process for making the data in the secondary logical
30 volume 72 capable of being accessed from the host computer 10 as a separated volume from the primary logical volume 71 will be

described. (This process is referred to as the "process for rendering the data accessible" hereinbelow). The "process for rendering the data accessible" may be incorporated in the OS 12 (or performed by the OS 12) or may be realized by executing an application program running on the OS 12. In the following description, it is assumed that the "process for rendering the data accessible" is enabled by executing an application program. Also in the following description, it is assumed that the host computer 10 is a mainframe, and the OS 12 is an operating system for the mainframe.

The "process for rendering the data accessible" includes:

a first process performed before making the control states of both volumes 71, 72 change to the "split state";

a second process for making both the primary and secondary logical volumes 71, 72 change to the "split state" and setting the management information so that the secondary logical volume 72 can be accessed as a separate volume independent of the primary logical volume 71; and

a third process for reconfiguring (re-setting) the management information so that the control states of both volumes 71, 72 can be changed from the "split state" back to the "paired state", and then actually making the control states of both the primary and secondary volumes 71, 72 change back from the "split state" to the "paired state".

--First process--

Referring to the flowchart shown in Fig. 2, the first process will be described. First, the application program obtains, from the user interface (not shown) of the host computer, a new volume name and a new data set name to be assigned to the secondary logical volume 72 that is to be rendered accessible independent of the

primary logical volume 71 (S211). These names can be obtained from information stored in the memory (not shown) of the host computer 10 or in the disk drives 25, instead of obtaining them from the user interface.

5 Then, the application program obtains the volume name of the primary logical volume 71 and the data set name for the data set stored in the primary logical volume 71 from the catalog information (S212). The catalog information is the management information of the OS 12 and includes such information as physical
10 addresses specifying the storage areas on the disk drives 25 and the attributes of the disk drives 25, which are associated with the data set name.

 The catalog information is stored in a predetermined storage area in the logical volume 70. The user manages the data
15 set in association with the catalog name. It should be noted that although in this embodiment the OS 12 is described to be one for a mainframe, if the OS 12 is for an open system, the catalog name may be the directory name, and the data set name may be the file name in a file system,.

20 Then, the application program loads dump information of VTOC stored in the primary logical volume 71 onto the memory or outputs the information to a data set (S213). The VTOC is management information managed for each logical volume and contains (1) the data set name, (2) the creation date, (3) the
25 method in which the file was organized, (4) the record format, (5) the location at which the data set is stored, etc. When accessing a data set on a logical volume, the OS searches for the storage location of the VTOC (VTOC address) from the volume label to retrieve the storage location (extent address) of the target
30 data set and uses the extent to access the data set.

 Next, the application program generates a first control

program (S214). The first control program is used to search, from the loaded or output dump list of VTOC, for the section in which the new volume name and the new data set name that have been obtained in step S211 are described, and then perform processes
5 for setting the thus-found volume name and the data set name in the VTOC of the secondary logical volume 72 to the volume name and the data set name obtained in step S212.

Then, the application program loads onto the memory, or outputs into a data set, the dump information of VVDS stored in
10 the primary logical volume 71 (S215). The VVDS is the information on the VSAM data set as described above and data sets subject to system management.

Then, the application program generates a second control program (S216). The second control program is used to search,
15 from the loaded or output dump information of VVDS, for the section in which the new volume name and the new data set name that have been obtained in step S211 are described, and then perform processes for setting the thus-found volume name and the data set name in the VVDS to the volume name and the data set name obtained
20 in step S212.

Then, the application program obtains the catalog name in which is recorded the secondary logical volume 72 to be made accessible from the host computer 10 as a volume independent from the primary logical volume 71 (S217). The catalog name may be
25 obtained from the information stored in the memory or the disk drives 25, instead of obtaining it from the user interface. If the catalog name thus obtained here is a new catalog name (S218: YES), the application program generates a third control program for recording in the catalog information a new catalog
30 corresponding to that catalog name, and a fourth control program for deleting the catalog from the catalog information (S219).

Then, the application program generates a fifth control program for recording the data set having the data set name obtained in step S212 in the catalog having the catalog name obtained in step S217, and a sixth control program for deleting
5 the data set (S220).

Finally, the application program stores the first through sixth control programs generated as above into the memory of the host computer 10 or a given data set (S221).

10 --Second process--

Now referring to the flowchart shown in Fig. 3, a second process will be described. At first, the application program sends a pair split command to the disk array device 20 to make the control state of the primary logical volume 71 and the
15 secondary logical volume 72 change from the "paired state" to the "split state" (S311). Then, the application program sets the volume name (volume label) of the secondary logical volume 72 (S312) to the name newly obtained in step S212. The secondary logical volume 72 is set in an online state with the new volume
20 name (S313).

Then the application program executes the first control program generated in step S214 to set the volume name and the data set name described in the VTOC of the secondary logical volume 72 to the new volume name and the new data set name (S314). In
25 addition, it executes the second control program generated in step S216 to set the volume name and the data set name described in the VVDS of the secondary logical volume 72 to the new volume name and the new data set name (S315).

Then, if the application program is to record the secondary
30 logical volume 72 in a new catalog (S316: YES), then it executes the third control program generated in step S219 to generate a

new catalog with the new catalog name obtained in step S217. Then, the application program executes the fifth control program generated in step (S220) to record the data set of the secondary logical volume 72 with the new data set name in the catalog with the new catalog name obtained in step S217 (S318).

Fig. 4 is a schematic diagram illustrating the states of the volume name, the data set name, and the catalog name when the secondary logical volume 72 is recorded in a catalog (CATALOG2) different from that (CATALOG1) of the primary logical volume 71. Fig. 5 is a schematic diagram illustrating the states of the volume name, the data set name, and the catalog name when the secondary logical volume 72 is recorded in the same catalog (CATALOG1) as that of the primary logical volume 71. Fig. 6 is a schematic diagram illustrating the states of the volume name, the data set name, and the catalog name when the secondary logical volume 72 is recorded both in the same catalogs (CATALOG1, CATALOG2) as those of the primary logical volume 71 and in catalogs (CATALOG3, CATALOG4) different from those of the primary logical volume 71.

Fig. 7 shows an example of a control program when the secondary logical volume 72 is to be recorded in a catalog different from that of the primary logical volume 71. Fig. 8 shows examples of the first, second, and fifth control programs when the secondary logical volume 72 is to be recorded in the same catalog as that of the primary logical volume 71. Fig. 9 shows an example of the first, second, fifth, and sixth control programs when the secondary logical volume 72 is to be recorded both in the same catalogs as those of the primary logical volume 71 and in catalogs different from those of the primary logical volume 71.

After the above processes, the OS 12 of the host computer will be able to have access to the secondary logical volume

72 as a logical volume independent of the primary logical volume 71. Thus, the OS 12 as well as programs of the host computer 10 will be able to use, for example, the data in the secondary logical volume 72 for backup and/or to perform any batch processing and
5 data analysis.

--Third process--

Now referring to the flowchart shown in Fig. 10, the third process will be described. At first, the application program will
10 execute the sixth control program to delete the data set with the data set name obtained in step S212 from the catalog with the catalog name obtained in step S217 (S811). Then, the application program sets the secondary logical volume 72 in an offline state (S812). Then the application program sends a resync
15 (resynchronization) command to the disk array device 20. In this way, the control state of the primary logical volume 71 and the secondary logical volume 72 is changed from the "split state" back to the "paired state" (S813).

As have been described above, the application program
20 automatically generates the control programs for automatically setting the volume name and dataset name described in such management information as VTOC and VVDS to a new volume or a new data set name, and for recording and deleting a new catalog. Thus, the user does not have to perform settings of the contents of
25 management information every time it is necessary in order to access the secondary logical volume 72 as a volume independent of the primary logical volume 71. As a result, the burden on the user to perform settings is significantly reduced.

Storing the first through sixth control programs thus
30 generated in the above-described processes in a given data set makes it possible to reuse the stored data set in situations where,

for example, operations such as regular data backup and daily batch processing are to be performed routinely or from time to time. This makes it possible to eliminate processes of setting the volume name and the data set name every time it is necessary to perform those jobs such as backup processing and batch processing. By making use of those control programs, it is easy to make the secondary logical volume 72 accessible from the OS 12 of the host computer 10 as a logical volume independent of the primary logical volume 71.

10 Although the control method of a storage system has been described in detail in accordance with an embodiment of the present invention, it should be understood that the foregoing description of the embodiment has been presented for purpose of facilitating understanding of the present invention and is not to limit the invention to the precise form disclosed. The present invention may be altered and/or modified without departing from the scope thereof, and it is without saying that equivalents of the present invention are included within its scope.

20 --Other embodiments--

 In the foregoing description, an embodiment having primary and secondary logical volumes 71, 72 managed by a single disk array device 20 has been presented. However, the present invention may be equally applicable to an embodiment in which a primary logical volume 71 and a secondary logical volume 72 are each provided by different disk array devices in separate housings. In such a case, the present invention can not only be applied to a case in which the disk array device 20 for providing the primary logical volume 71 and the disk array device 20 for providing the secondary logical volume 72 are placed close to each other such as in the same building, but also to a configuration in which the disk array

devices are remotely located away from each other and connected through a communications line.

In addition, it is also possible to provide a mechanism for making a notification to a user through a user interface used for
5 designating a volume name etc. to be set, in case the volume name or the data set name to be set is already used.

In brief, in accordance with the present invention, a novel control method of a storage system, a storage system, and a control program therefor may be provided.

10 Although the present invention has been described in conjunction with several preferred embodiments thereof, it should be understood that these embodiments are disclosed by way of examples and the present invention is not to be limited thereto. It should be recognized that many changes and modifications may
15 be made by those skilled in the art without departing from the true spirit and the scope of the present invention set forth in the appended claims.